



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-643863

Application of Subspace Detectors to 2012 Sumatera Earthquake Sequence

D. Dodge

September 16, 2013

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Application of Subspace Detectors to 2012 Sumatra Earthquake Sequence

Doug Dodge
Lawrence Livermore National Laboratory

Abstract

In this report the aftershock sequence of the 2012 Sumatra earthquake is used to study the performance of subspace detectors (Harris, 2006) to detect and classify the events. The ground truth used for this study is 1222 aftershock solutions from the LLNL database drawn mostly from the Reviewed Event Bulletin (REB). Subspace detectors are built using several spawning strategies with waveforms recorded by the Makanchi Array (MKAR).

Because we desire the subspace detectors to have a very low false alarm rate, the subspace templates are required to be about 50 seconds long (most of the P-bundle). For the same reason, the power detections used to create the subspace templates are restricted to have high SNR. Because of these restrictions, none of the subspace detectors are able to reproduce the entire ground truth catalog. They do have a very low false alarm rate, however, and so can reliably be used as single-array detection classifiers.

Because of their reliability, in a suitably designed system, these detectors could be used to segregate detections from regions of high seismicity such as aftershock sequences, and allow processing of those detections in isolation from routine seismicity. It may also be possible to use these detectors to process groups of detections simultaneously. This leads to the concept of “analyst workload reduction factor”; a measure of the reduction in effort achieved by processing N detections in M groups instead of individually. If M is much smaller than N , a significant reduction factor is achieved.

We use three subspace detector spawning strategies.

1. In the first strategy, correlators are spawned directly from power detections and are allowed to run to the end of the sequence. Correlation clustered detections are then used to create multi-rank subspace detectors which are run against the entire sequence.
2. In the second strategy, power detectors are run for the entire duration of the sequence. Nearly the entire set of detections is used to create one or more subspace detectors with a high energy capture value.
3. The final strategy attempted was to build subspace detectors as in (2) but restricted to only the first day's detections.

The second strategy was the most successful. Four subspace detectors produced 781 detections for the 44-day period with at most a few false detections. Operationally, this system is not practical, but we

present a method for creating the detectors incrementally which we believe would maintain a high sensitivity with few false detections, while being able to provide detections without a delay for template creation.

Introduction

On April 11, 2012 a pair of large (M8.6, M8.2) strike-slip earthquakes occurred off the west coast of Northern Sumatra, Indonesia. The earthquakes were accompanied by a strong aftershock sequence with ~1220 events recorded in a 44-day period following the first event of the sequence. Of these about 70 were of magnitude 5 or greater.

Earthquake sequences like this are a problem for monitoring agencies such as the International Data Center (IDC) because the high rate of activity can make it difficult for analysts to keep up with the evolving sequence. Consequently, there is considerable interest in developing processing techniques for detection and association that are more robust in these conditions.

Bobrov et al., 2012 have used cross correlation as an aid in analyzing aftershocks of this sequence. They processed 44 days of data from 7 IMS arrays using as templates, waveforms from 16 master events chosen to cover the spatial extent of the aftershock sequence. Their processing produced about 1,000,000 cross correlation detections. Many of these detections were removed in a conflict resolution process so that in the end, they produced 4924 event hypotheses for the 44-day period. In all, 2763 new event hypotheses (relative to the SEL3) were created.

In 2013, the IDC hosted a group of scientists from three US Department of Energy laboratories for a one-week series of meetings to discuss correlation processing of aftershock sequences. During this meeting, the DOE scientists were briefed on the IDC's correlation processing methodology as applied to the Sumatra sequence. The DOE scientists also agreed to research various aspects of this problem in hopes that the combined contributions might lead to further advances in applying correlation processing to earthquake monitoring. The representatives from Lawrence Livermore National Laboratory (LLNL) agreed to process this sequence using a subspace detection framework which they have been developing for several years. This report summarizes the results of that work.

The Detection Framework

The subspace detection framework used for this work was first described in Harris and Dodge, 2011. The system processes array data and acquires waveform templates with an STA/LTA detector operating on a beam directed at the P phases of the aftershock sequence. The templates are used to create correlation-type (subspace) detectors that sweep the subsequent data stream for occurrences of the same waveform pattern. Events are clustered by association with a particular detector.

The assumption in developing this system was that in many (if not most) sequences a significant fraction of aftershocks cluster around asperities. If these clusters are spatially compact and have similar mechanisms, then it should be possible to use correlation to identify the clusters and use

that information to guide association and even to provide analysts with cluster-level views of the data. A pipeline that implements such a strategy could in principle be able to reduce analyst workload by $(N-C)/N * 100\%$ where N is the number of events in a sequence and C is the number of clusters.

The first application of the framework was to process 10 days of aftershocks of the 2003 San Simeon earthquake recorded by the NVAR array approximately 390 km away from the source region. In that application, the detection framework operated in a 2-pass mode produced 676 detections using 184 detectors. The theoretical workload reduction was nearly 73%.

Since the initial implementation of the framework we have continued to develop the system under BAA funding (Kvaerna et al. 2011, 2012). In addition to a number of performance improvements, the system now includes fairly sophisticated screening mechanisms to prevent creation of templates from unwanted power detections. The system also includes the ability to create and use matched field detectors, to perform multiple passes over the data, and to perform detector maintenance automatically.

The Data

Following Bobrov et al., 2012, we analyzed 44 days of activity in a region extending in latitude from -1.5° to 6° and in longitude from 86° to 96° . The LLNL database contained 1222 events for those constraints (1219 from the IDC REB and the remainder from the PDE daily bulletin). Figure 1 shows the distribution of events in the upper panel. The location of the main shock is marked with a red cross. These events serve as the ground truth for evaluating the performance of the detection framework. We had initially hoped to use detections from station PSI (Parapat, Indonesia) as ground truth as it is only a few degrees from all the aftershocks. Unfortunately because of high background noise we did not get enough detections there to make it useful for that purpose.

The lower panel of Figure 1 shows the event rate as a function of time. Note that by about the tenth day the rate of aftershock production has declined to the point where the utility of automated aftershock processing is becoming marginal. A system which has not helped with processing the events up to that point is probably not very useful as part of a processing pipeline. We make this point because, of the two most successful experiments with the detection framework, one required 44 days' worth of power detections to create its templates and one required a full day's worth of power detections. However, we believe that by employing SVD updates to create the eventual high-rank subspace templates, this time lag before the templates are useful could be eliminated.

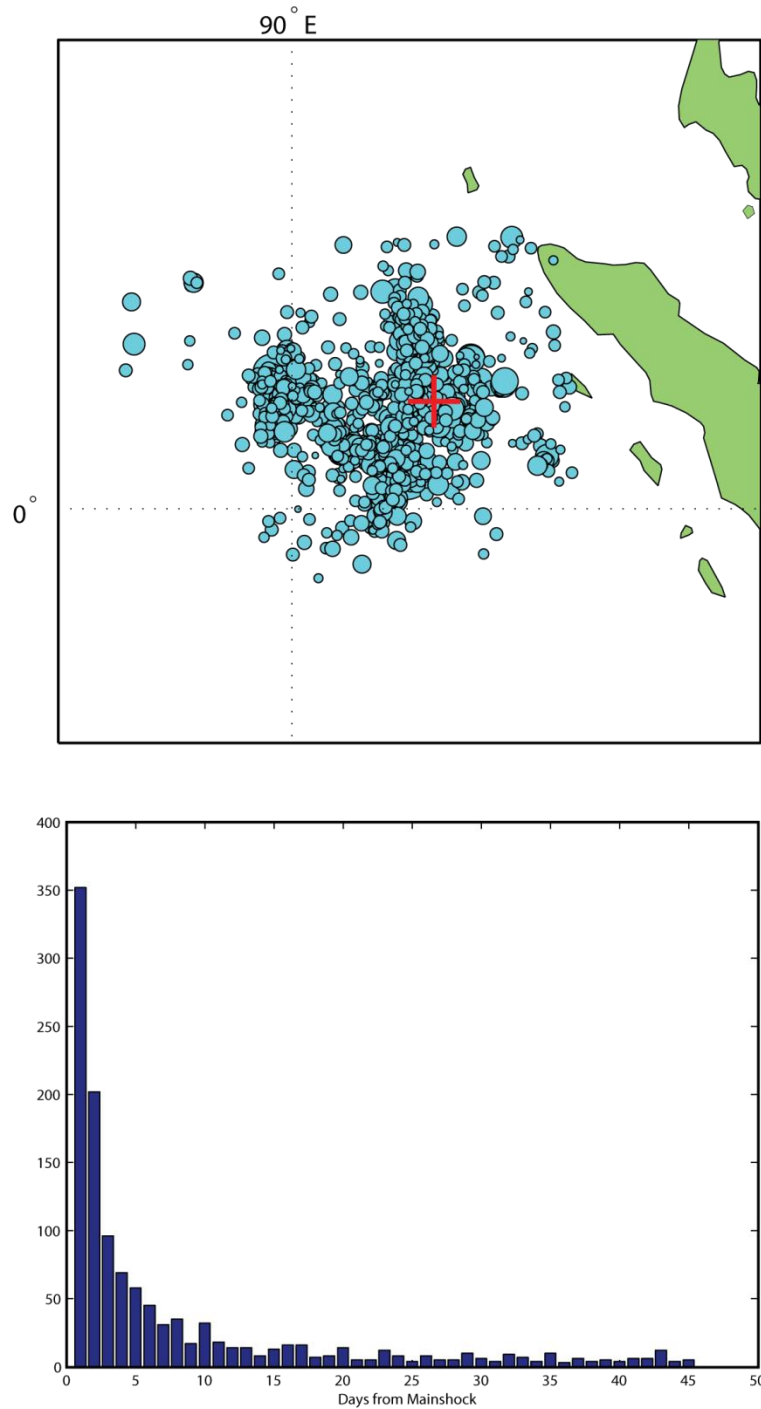


Figure 1 (Upper) The 1222 events from the LLNL database in the source region for 44 days following the main shock. (Lower) Events per day for this dataset. Note that by day ten only a few tens of events are being produced in the source region.

Our original intent was to run the detection framework on data from the same seven IMS arrays as was used by Bobrov et al, 2012. These are MKAR, WRA, CMAR, SONM, ASAR, ZALV, and GERES. In fact we did some processing on 5 of these arrays (ASAR, CMAR, GERES, MKAR, and SONM). Figure 2 shows detections as a function of time for these five arrays.

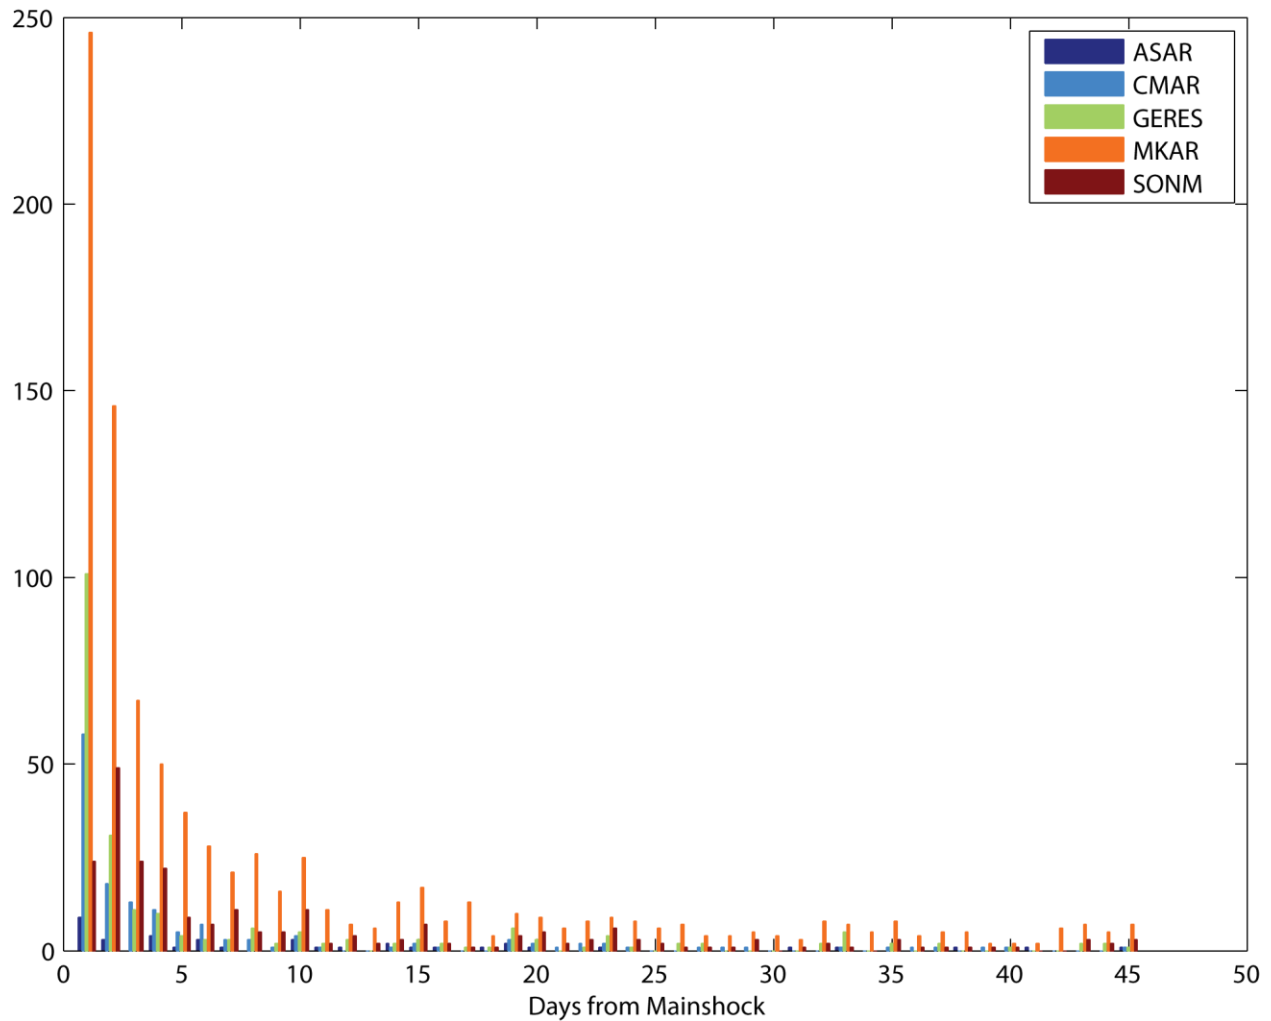


Figure 2 Subspace detections as a function of time for 5 IMS arrays for 44 days after the 2012 Sumatra main shock.

Each configuration used in this test had four spawning detectors based on 4 events in the sequence that spanned the aftershock area. Each beam was configured for its array by performing an FK analysis. All subspace detectors used a threshold of 0.3. MKAR was the clear winner in this comparison. No other array outperformed MKAR on any day, and no other array detected any (catalog) event that was not also detected at MKAR.

Although it is possible to build configurations for the detection framework that employ multiple arrays, the resulting templates tend to be quite long if the arrays are at very different distances from the source region. This impacts the performance because of the increase in length of the FFT's that must be computed. In addition, such templates tend to contain a lot of noise because for each array it is necessary to include the time span where signal is expected at other arrays in the template. We are working on a solution to that problem, but for this study it would have been necessary to build long templates with very little signal. Given these considerations, we completed this study using only data from MKAR.

Choosing an Optimal Template Length

Prior to running the framework on the MKAR data it was necessary to decide on a template length for the subspace detectors to be generated. Bobrov et al., 2012 used window lengths of 4.5s to 11s, but as previously remarked, their system produced a large number of false detections that had to be removed by post-processing of the data. The detection framework is meant to be operated in such a way that few or no false detections are produced. Therefore, one must ensure that a sufficiently great time-bandwidth-product is achieved.

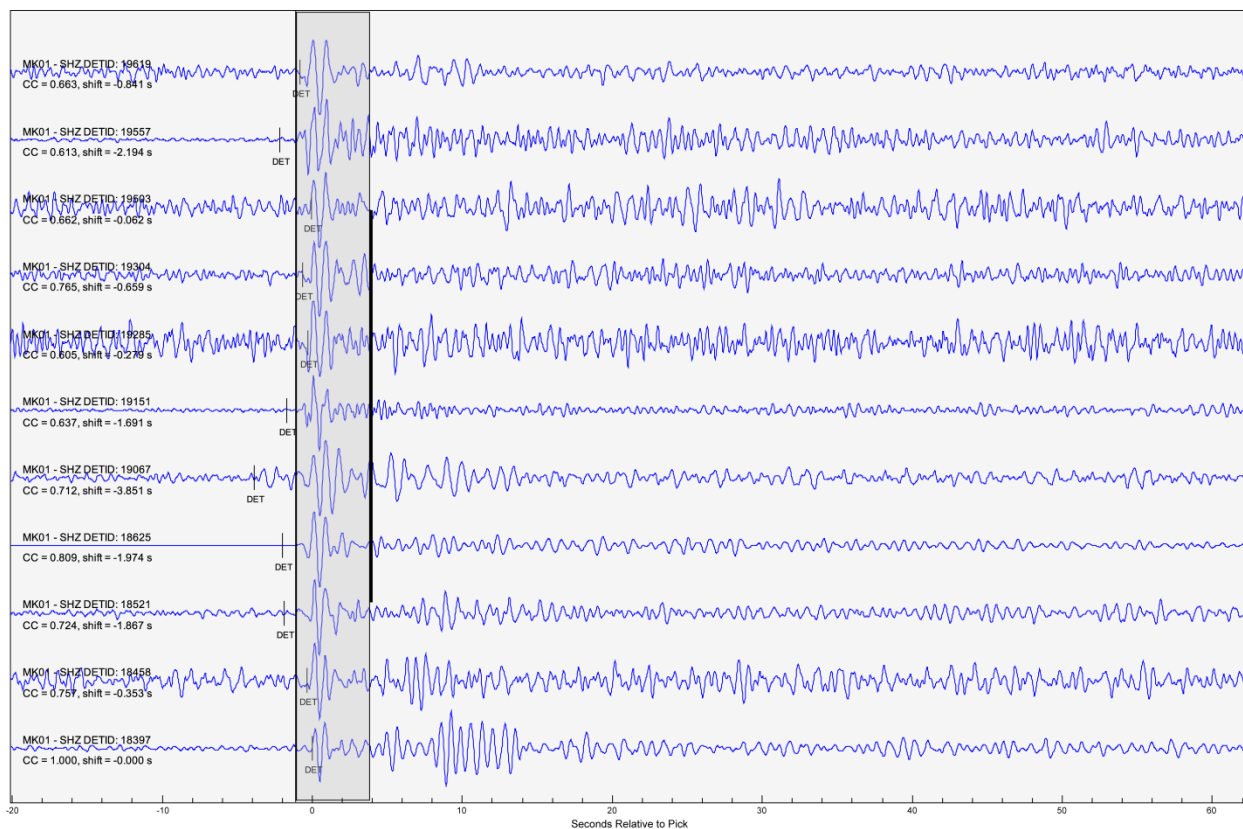


Figure 3 Example detections made using a 5-second long template. Note the wide disparity in the P-coda despite the similarity of the data in the template window.

Figure 3 above shows a few signals recorded at MKAR and correlated using 5-second windows in the 1-3Hz band. These short windows capture only a few cycles of the P-wave. Although within the window, the signals are similar, there is a wide variation in the rest of the signal. This is not just a matter of interfering signal. For example signals 2, 6, 8, and 11 all have high SNR, yet the coda are quite different. These events are unlikely to be similar in both position and mechanism. It is clear that longer windows are required to avoid misclassifying these signals. However, the longer the window, the greater the probability of the signals becoming de-correlated because of interfering signals.

To determine an optimal length we analyzed a set of 419 power detections at MKAR. This analysis assumes that the P-coda is mostly composed of scattered energy from the events being measured, and that there must be a window length beyond which little additional information about the Green's function and source is present. Using windows longer than that will increase the risk of contaminating the signal while providing little improvement in the information content.

Table 1 Largest cluster size as a function of window length and cluster threshold.

Window length (s)	5	10	20	30	40	50	60	70	80
Cluster size at .6	208	112	43	26	16	6	4	3	6
Cluster size at .7	125	29	7	4	3	3	3	3	3
Cluster size at .8	29	8	3	3	3	3	3	3	3

For nine different window lengths we cross correlated and clustered the waveforms using (squared CC) thresholds of 0.6, 0.7, and 0.8. The metric used was the size of the largest cluster. The results are shown in Table 1 and in Figure 4.

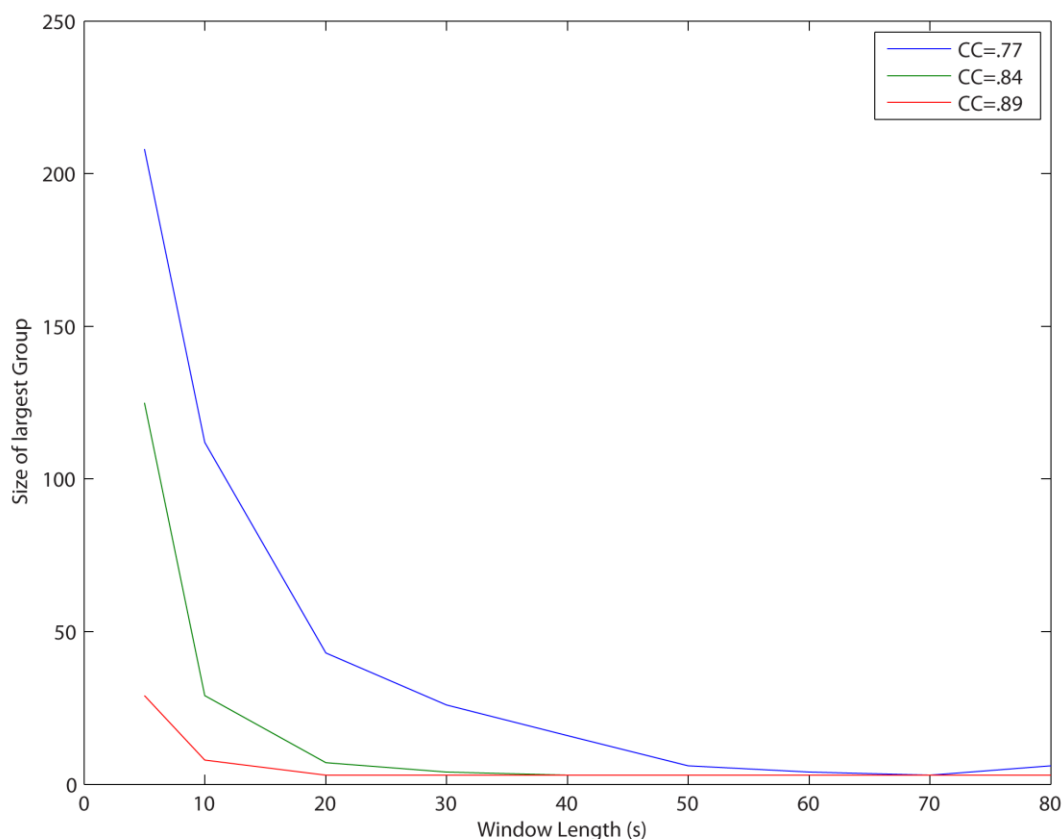


Figure 4 Largest cluster size as a function of window length and cluster threshold shown graphically.

Evidently, at the higher thresholds, a window length of 20s – 30s appears to be long enough to avoid spurious correlations, but at a threshold of 0.6 (.77) about 50 seconds is required. Because we seek to maximize the number of correlation detections, we must use a low threshold, and so choose the length of 50s.

Determination of Screening Parameters

After deciding on a template length to use at MKAR we performed a series of runs to investigate how the framework performs on teleseismic data. (Previously, we have only analyzed regional data.) These runs also were used to establish the screens that are applied to power detections before they are used to create templates. This screening of detections is very important in order to avoid creation of templates that are sensitive both to the intended signal and to unwanted signal or noise.

Four power detectors were used as spawners for these experiments. The detector beam parameters were chosen by selecting 4 aftershocks from the LLNL catalog whose positions spanned the sequence and obtaining the azimuth and velocity at MKAR for these events by FK analysis. These are shown in Table 2.

Table 2 The beam parameters for four power detectors used for spawning of correlators.

LLNL EVID	Back Azimuth	Velocity	Detectorid
6818030	166.6	11.8	10419
6872645	172.3	13.1	10420
6899681	163.3	13.1	10421
6900130	163.3	11.9	10422

To establish the bounds for FK screening, we performed a run of the framework with the configuration set to only write out power detections and with no screening of triggers by FK. The upper-left panel of Figure 5 shows the FK results for those detections. Despite the beam-forming being used by the power detectors, there are numerous detections from all azimuths and a wide range of velocities. By identifying which detections were in the catalog, we were able to determine that bounds of $\pm 15^\circ$ in azimuth and $\pm 1.k$ km/s were appropriate.

After determining the appropriate bounds and removing detections that did not meet the screening criteria, we created 18 subspace detectors from clusters of similar detections. These detectors plus correlators produced by the framework were run at 5 different threshold levels to see how detection performance behaved as the thresholds were decreased. The results are shown in Table 3 below.

Table 3 Detection and detector counts for the five framework runs.

Run ID	Threshold	Total Detections	Contributing Detectors	Workload reduction
101	0.5 (.707)	998	918	8%
103	0.4 (.632)	1179	1041	12%
104	0.3 (.548)	1284	1020	21%
105	0.2 (.447)	1691	970	43%
106	0.1 (.316)	5276	940	82%

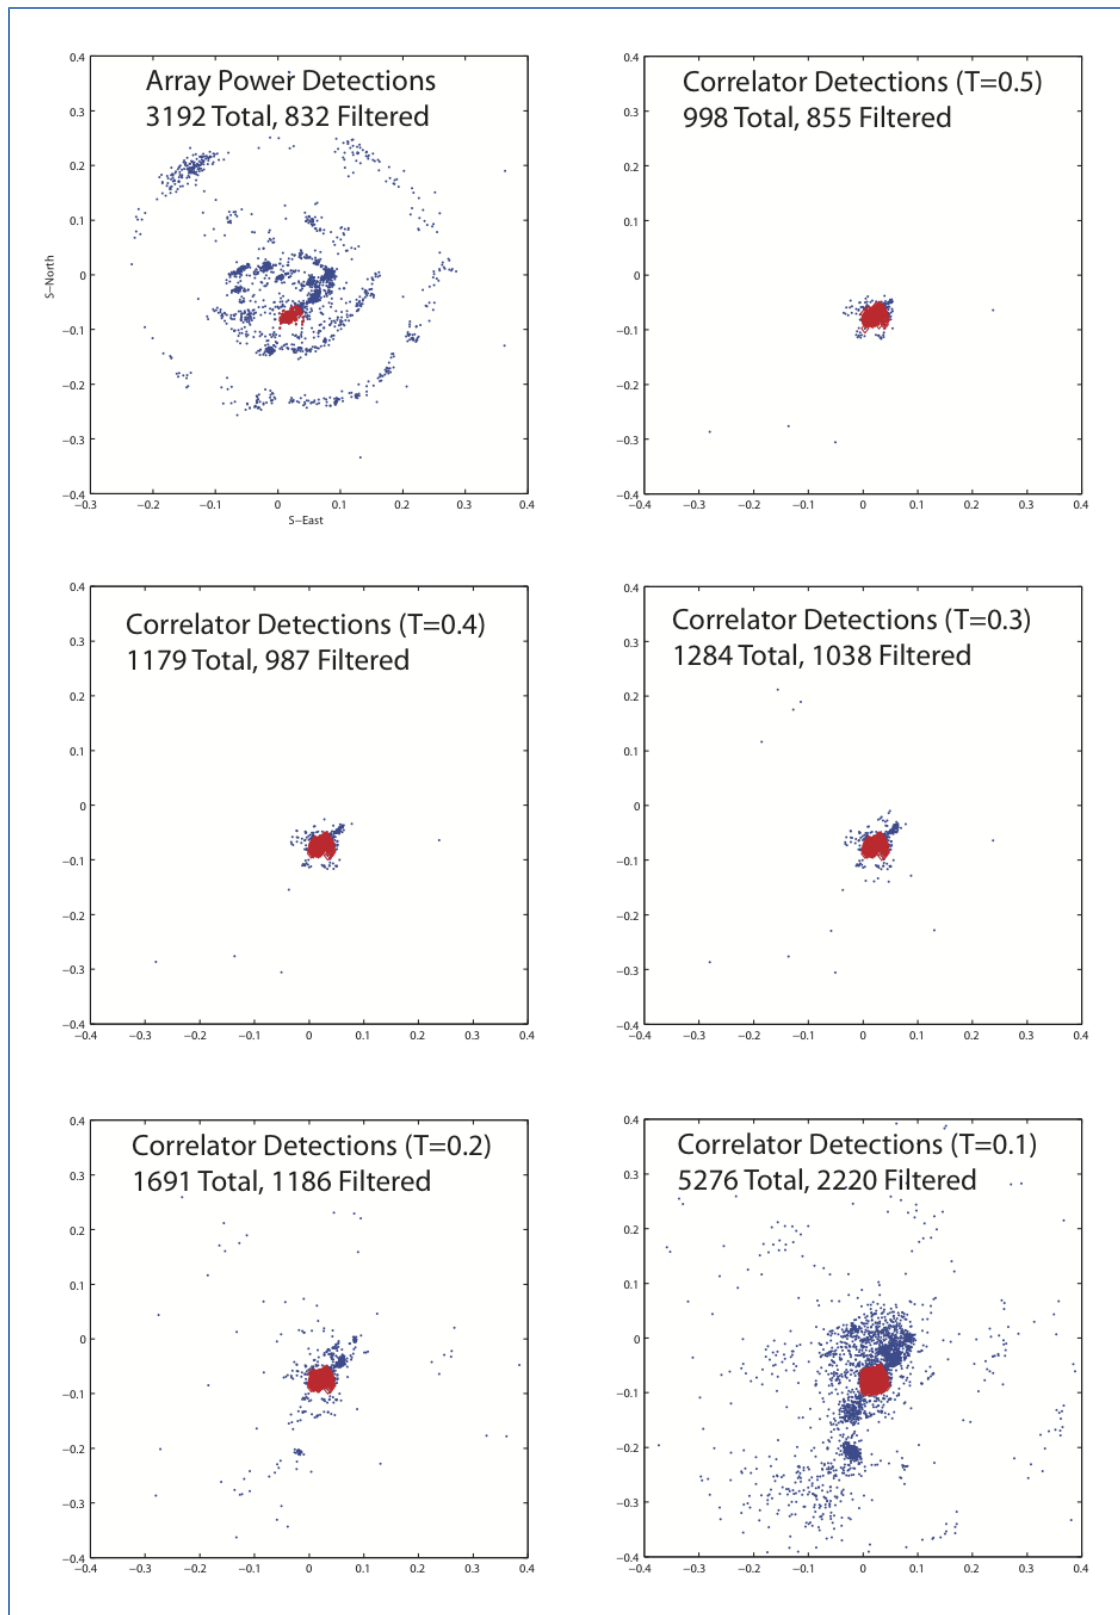


Figure 5 FK results for the power detections and for the correlators at 5 values of T . Upper left panel shows results for the power detector. Note the large numbers of side-lobe detections.

Note that although the framework has succeeded in producing more detections in every run than could be found using a power detector, in every case, about 1000 detectors are required. In other words, very few of the correlators are detecting anything other than their template event. The signal on which their template is based is unique.

Not surprisingly, the detection count increases with decreasing threshold. However, below a threshold value of 0.3 the detection counts appear to be suspiciously high. To investigate, we extracted the FK results for all the detections. They are plotted in Figure 5 along with the FK results from the power detector. Evidently, there are a large number of false detections. Even at $T = 0.5$, some of the detections have slowness that diverges greatly from that of their pattern event. At $T = 0.1$ performance is abysmal. Most of the detections appear to be false.

Our first thought was that the multi-rank subspace detectors were adapting themselves to data in such a way as to detect signals with very different slowness values. We extracted the detection counts (filtered by FK and unfiltered) by detector and show some selected values in Table 4. The fraction column reports the fraction of total detections that also passed the FK screen. The multi-rank detectors did as well as or better than the rank 1 detectors in detecting signals with the correct slowness.

Table 4 Selected detectors with their detection counts from run 106. Both the unfiltered and filtered counts are shown as well as the ratio of filtered to unfiltered. Multi-rank detectors are shown with yellow shading. The two worst-performing detectors are shown with red shading.

Detectorid	Unfiltered count	Filtered count	Fraction
6029	1545	914	0.591
6803	715	10	0.014
5977	639	1	0.001
5951	183	117	0.639
6823	162	42	0.259
6123	110	78	0.709
6667	103	17	0.165
6465	95	2	0.021
5961	88	43	0.489
6727	64	25	0.391
6351	54	2	0.037
5966	46	33	0.717
6314	35	33	0.943

Table 4 also shows some detectors that performed extremely poorly. For example only about 1% of the detections from detector 6803 met the FK screen. Detector 5977 was even worse with only about 0.1% passing the screen.

We extracted the templates for detectors 6029, 6803, and 5977, and have plotted channel MK01 for each of these templates in Figure 6. Clearly, an important contributor to the large number of false detections in run 106 was the creation of templates from power detections with low SNR. The top two templates in Figure 6 have SNR of 1.9 and 2.9 respectively, and between them contributed 1343 of the 3056 detections that failed the FK screen in run 106. By contrast detector 6029 ($SNR = 90$) had 59% of its detections pass the FK screen.

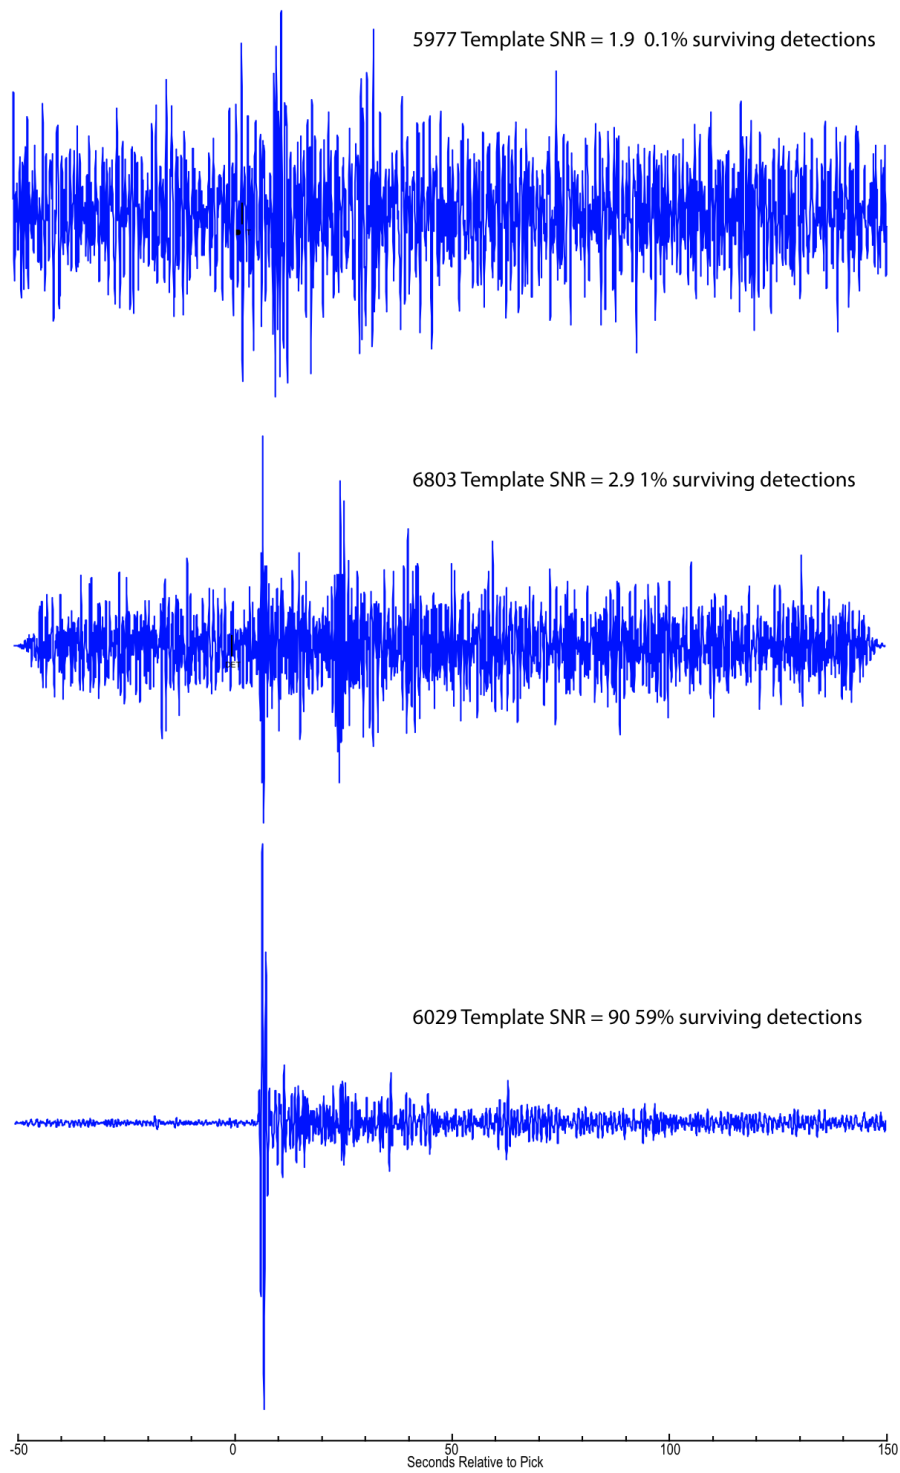


Figure 6 Channel MK01 from three templates arranged in order of increasing SNR. The top two built large numbers of bogus detections, while the bottom one with SNR = 90 performed fairly well at $T = 0.1$.

Subspace Detections From Correlation-derived Clusters

With screening parameters determined, we were ready to see how well the framework could do in reproducing and extending the catalog of detections using our technique of building subspace detectors from correlation-derived clusters. The process proceeds in 3 steps:

1. Run the framework in the mode that produces correlation detectors from power detections.
2. Build multi-rank subspace detectors from the detections of the best-performing correlators.
3. Run the framework with the subspace detectors in place.

Step 1 produced 366 detections and generated 209 detectors. The top 10 detectors had 131 detections among them as shown in Table 5.

Table 5 The top ten performing detectors from first run of the framework using automatic spawning of correlators and a screening threshold of 5.0.

Detectorid	Detection Count
11156	32
11258	24
11150	20
11164	16
11149	9
11210	7
11253	6
11218	6
11159	6
11225	5

Using detections of the top seven detectors, we created the multi-rank subspace detectors shown in Table 6. Running with this augmented suite of detectors (209 correlators and 7 subspace detectors) the framework produced 480 detections using 215 detectors. Table 6 lists the seven subspace detectors along with their rank and the number of detections they produced at the 0.3 threshold.

Table 6 Multi-rank subspace detectors created from detections from first run of the framework.

SS Detectorid	Source Detectorid	Rank	Detections in follow up run
11356	11156	18	129
11357	11258	5	20
11358	11150	7	29
11359	11164	11	43
11360	11149	5	11
11361	11210	4	6
11362	11253	4	7

How well Did we Do?

Figure 7 shows the detections produced by the second run of the framework. The white filled circles show the event locations from the bulletin. The blue circles represent detections by correlators, and the red circles represent detections from multi-rank subspace detectors. The pattern detectors were effective mostly in the center part of the sequence where the power detectors had most of their detections. The pattern detectors out-performed the power detectors by 31% but only detected 39% of the events in the catalog.

From the standpoint of being a highly-sensitive detector, this configuration was greatly outperformed by the system described by Bobrov et al, 2012. Recall that their system produced 4924 event hypotheses (approximately 4 times the number of events in our catalog and more than 10 times the number of produced by the framework). However, their system works by producing very large numbers of detections at many arrays, a large number of which must be removed by statistical analysis.

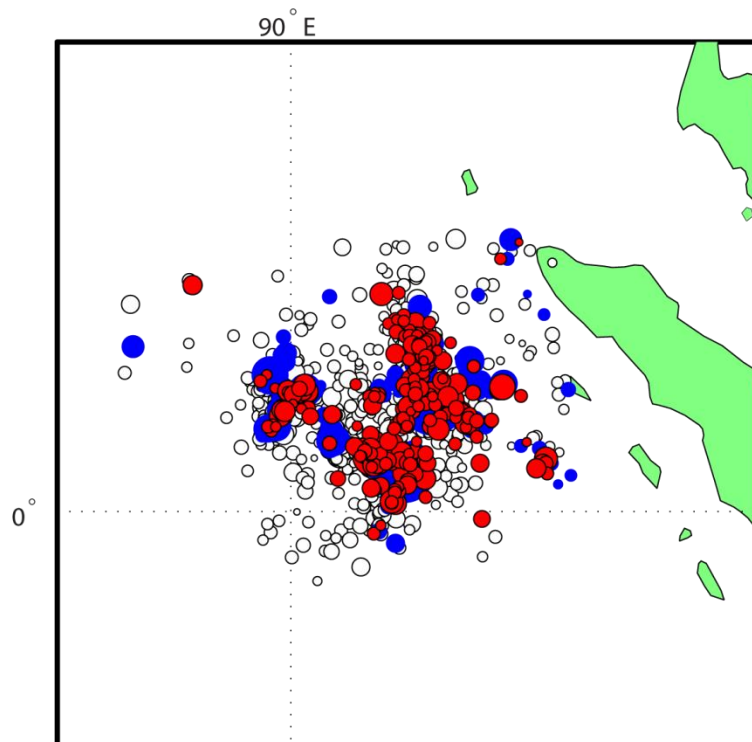


Figure 7 (White) Events from bulletin. (Blue) Correlation Detections. (Red) Subspace Detections.

In contrast, the pattern detectors presented here use data recorded at a single array to produce a set of detections that can be reliably assigned to a very compact source region. As part of a tool chain that allows groups of detections to be processed together there is the potential for significantly easing system load during an intense aftershock sequence. At a minimum the system could de-queue detections known to be from an aftershock sequence so that events from other parts of the world could be processed in a timely manner. More sophisticated strategies could be developed to allow joint processing of groups of aftershocks.

Given 480 detections using 215 detectors, this run demonstrates a potential analyst workload reduction of 55%. These results are not as good as what we achieved with the San Simeon sequence (73%). But, considering just the subspace detectors, performance is much better. There were 245 detections by 7 subspace detectors yielding a workload reduction of 97%. The limitation of the multi-rank subspace detectors is the small size of their footprint. In the next section, we present a way to substantially increase the size of the footprint.

Increasing the Footprint Using Very High-Rank Subspace Detectors

The subspace detectors used in the previous section were built using waveforms from events that had been clustered by correlation. In other words, the basis seismograms for these subspace detectors is from a spatial region small enough that the seismograms were correlated with one another at least at the 0.3 level. So it is not surprising that the clusters produced by these detectors are spatially compact as well. We expect that these detectors can be used to find additional events from within their basis region, and they do. But we are unable to produce by correlation a small set of clusters that spans the entire aftershock region for the sequence under study.

However, there is no reason that the basis used to form a subspace detector has to be chosen by correlation. As pointed out by Harris (2006) subspace detectors can span the range between pure correlators (rank 1) and energy detectors (rank equal to the signal space). By forming detectors that use more of the signal space in their basis, we can increase the footprint at the expense of an increase in the false alarm rate.

This approach has been advocated by Harley Benz of the US National Earthquake Information Center where they are having success with an experimental system that employs subspace detectors formed using all power detections from a certain source region. Encouraged by Harley's reports, we decided to try the same approach with the Sumatra sequence.

We started with the set of power detectors shown in Table 2 and ran the framework for the 44 days following the main shock. Although we retained the FK screen and duration screen used in the previous section, we lowered the detection threshold from 5 to 0.5. This adjustment was motivated by a desire to increase the size of the basis used to construct subspace detectors.

This configuration produced 3953 detections that passed the FK and duration screens. However, inspection of the detections showed a number with very low SNR. Concerned about the effect these signals might have on the false alarm rate of the subspace detectors, we applied an additional SNR screen passing only those detections with $\text{SNR} \geq 3.0$. This left the 485 power detections illustrated in Figure 8.

Only about 40% of the catalog events are included in the power detections. Also, the power detections do not span the area covered by catalog events. As will be seen, this limits the footprint of the subspace detectors.

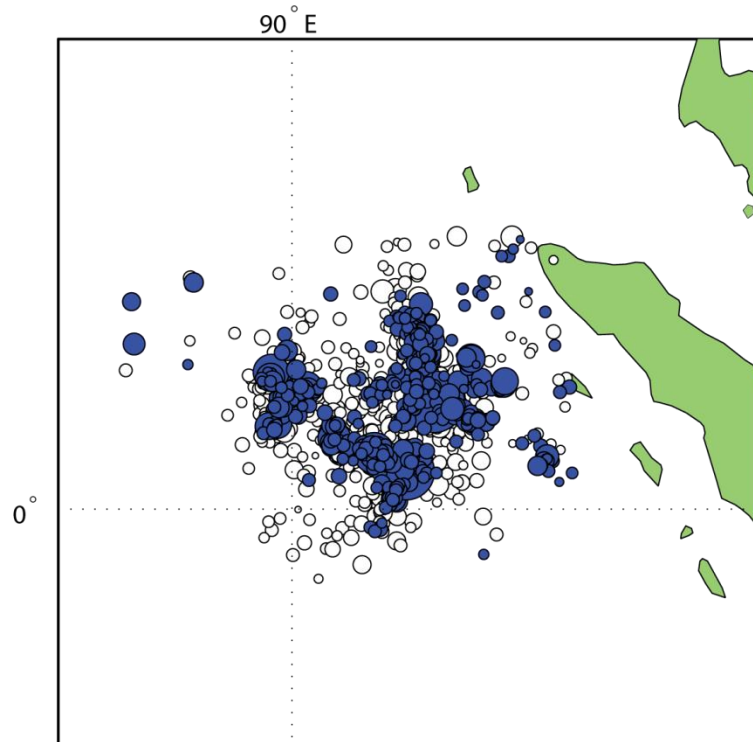


Figure 8 Power detections at MKAR (blue) plotted on top of the catalog locations (white). There are 1222 catalog events and 485 detections associated by time to catalog events. For the most part, the power detections are concentrated in the central part of the sequence.

For each of the four detectors, we extracted waveforms for all detections and used them to build subspace detectors using a template length of 50 seconds and an energy capture of 0.9. Table 7 lists the four detectors created this way along with their rank.

Table 7 Array Power Detections and Resulting subspace detectors.

Power Detectorid	Count	SS Detectorid	SS Detector Rank
10419	209	10426	115
10420	92	10425	78
10421	83	10424	71
10422	97	10423	56

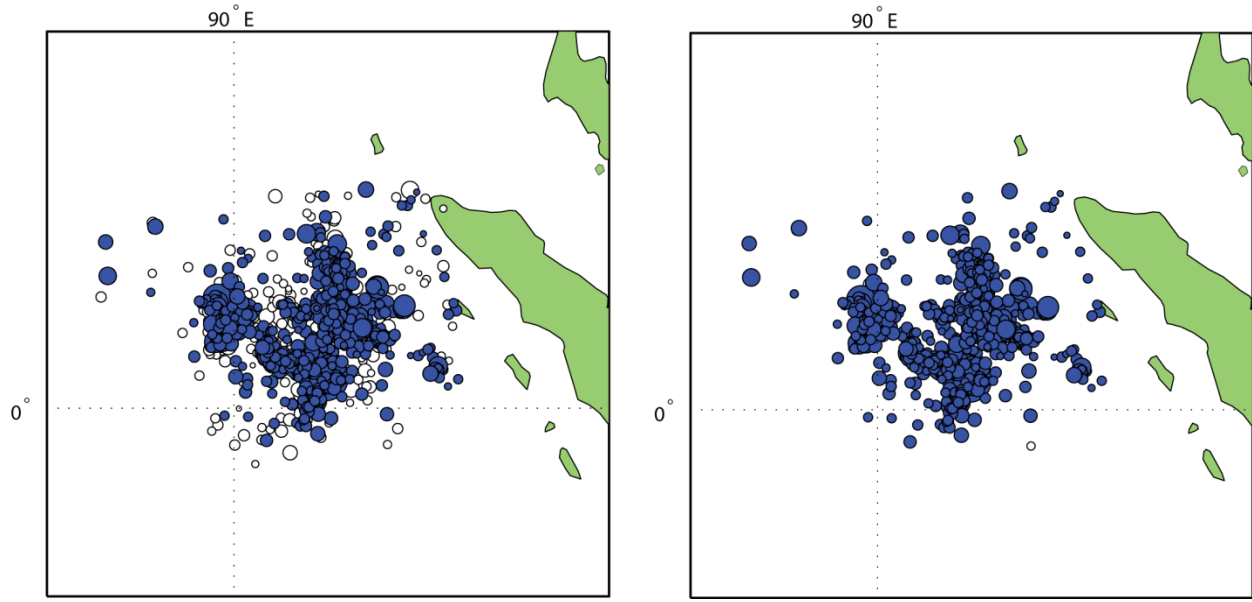


Figure 9 Detections for the 44-day period made using high-rank subspace detectors. (Left) Associated detections (blue) plotted on catalog events (white). (Right) Associated detections plotted on power detections (white).

Figure 9 shows the results obtained running the four high-rank subspace detectors against the 44 days of data. There were 781 subspace detections produced, 629 of which were associated to catalog events by time. In terms of absolute numbers of detections, this is a substantial (63%) improvement over our previous attempt using correlation-derived templates. Comparing Figure 9 (Left) to Figure 7 It seems that although most of the detections are still restricted to the central part of the sequence, within that region, there are significantly fewer missed detections (relative to the catalog) than was the case with the previous detector. Specifically, within a 4° by 4° box centered on $\text{lat} = 2^\circ$, $\text{lon} = 92^\circ$ the catalog contains 889 events while these detectors produced 567 associated detections (63%).

Where this configuration really shines, however, is in the analyst workload reduction factor. Assuming that detections can be dismissed or at least have their processing done as part of a group once the proper group has been determined, these detectors achieved a workload reduction of 99% $[(781-4)/781 * 100\%]$.

False Detections

The high-rank subspace detectors produced 152 detections that could not be matched to catalog events, and it is natural to wonder how many of those detections are false. All of the detections passed the FK screen with a minimum FK quality of 0.6, so we can at least say that there was coherent energy at the right slowness to be from the source region of the Sumatra aftershocks. However, for 28 of the detections, visual inspection of a single channel filtered into the frequency band of the detectors, did not reveal any obvious P-arrival. Some of those could be false detections.

We also ran the detectors over ten days of data prior to the main shock. This resulted in 12 detections. Five of those detections could be associated by time to events in the LLNL catalog with epicenters in the

aftershock zone. The remainder, although unassociated, are real events with appropriate slowness and good SNR. Figure 10 shows the detections on the MK01 channel.

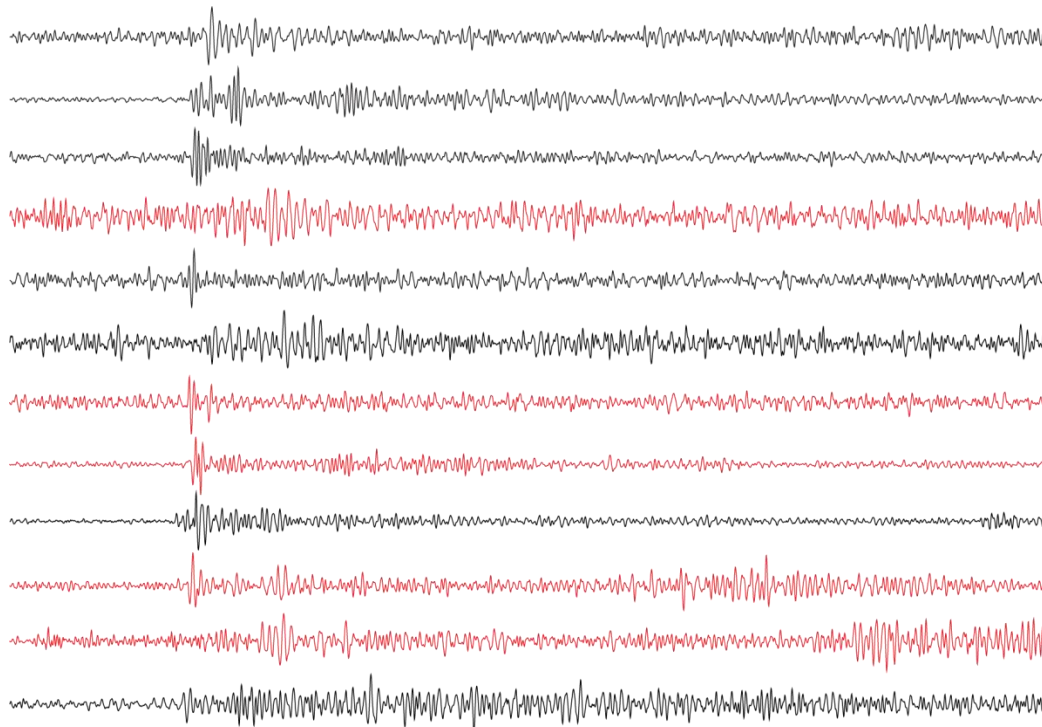


Figure 10 All detections (shown on MK01) from the 10 days preceding the main shock. The traces in red are for detections with times that associate with events in LLNL database. The associated events are in the aftershock zone.

The high-rank subspace detectors presented in this section have good performance, particularly at classification but have the disadvantage of having required the entire suite of power detections for the 44-day period in order to create the templates. It is no help to get these detectors only after the sequence has been fully processed. However, there exists a way to create the subspace basis incrementally, and although it has not been implemented for this report, it may perform comparably to (or possibly better than) the detectors just discussed.

Subspace templates are constructed by building a matrix whose columns are the basis waveforms, performing a Singular Value Decomposition (SVD), and retaining a subset of singular vectors as the template. However it is also possible to build up a SVD by adding a column at a time to the basis and performing a SVD update (Bunch and Nielsen, 1978) with each addition. Under this strategy, the first screened power detection would be used to form a correlator, and with each successive detection by either the power detector or the evolving subspace detector, the template would be updated. Initially, the analyst workload reduction factor would be 0. But, with each additional detection, the factor would increase. Because the template would include both power detections and subspace detections, it might turn out to be a more sensitive detector than one constructed using only power detections.

High-Rank Subspace Detectors Formed with a Subset of the Power Detections

An alternative strategy for bringing a high-rank subspace detector into service early in the sequence is to use as a basis only those events that occur early in the sequence. In this section we describe the performance of such a detector and compare it to the detector created using 44-days of detections.

The detectors described here were created using only detections from the first day of the sequence. Ten spawning detectors were used and ten subspace detectors were created from their detections. The intent in using a large number of spawners distributed around the aftershock zone was to minimize the footprint of each subspace detector.

Table 8 shows the statistics for the 10 spawning detectors including their detection counts. Note that 4 of the detectors had fewer than 10 detections at the end of the first day. In retrospect, those should probably have been discarded, but we used them anyway. The detectors that resulted from them were low-rank and contributed very little to the overall total.

Table 8 Statistics for the 10 spawning detectors including their detection counts after one day.

LLNL EVID	Lat	Lon	Mw	detectorid	Azimuth	Slowness	Detection count
6818018	1.5255	90.8811	6.06	10428	170.4	12.7	10
6818058	2.9313	89.5575	6.58	10427	170.7	12.3	4
6818086	3.2731	93.8179	5.14	10429	165.8	13.3	2
6818087	2.3937	93.451	5.25	10430	163.3	12.4	15
6818114	3.6929	92.7048	5.4	10431	167.2	13.3	10
6818166	4.372	92.7539	5.37	10432	162.8	13.6	19
6818171	1.2826	91.8917	5.37	10433	170.6	11.9	33
6818213	2.4692	92.7499	5.11	10434	166.8	12.4	3
6872788	1.8497	92.9442	5.51	10435	167.7	11.9	19
6986914	2.2752	91.7426	4.75	14036	167.4	12.7	3

Figure 11 (Left) shows the first-day detections from the 10 spawning detectors in blue and the events from the catalog for the entire 44-day period in white. The locations of the events on which the spawners were based are shown in red. The main point of this plot is to show that by the end of the first day the sequence was still rather compact, so the spawning detectors and their detections only sample the central part of the sequence.

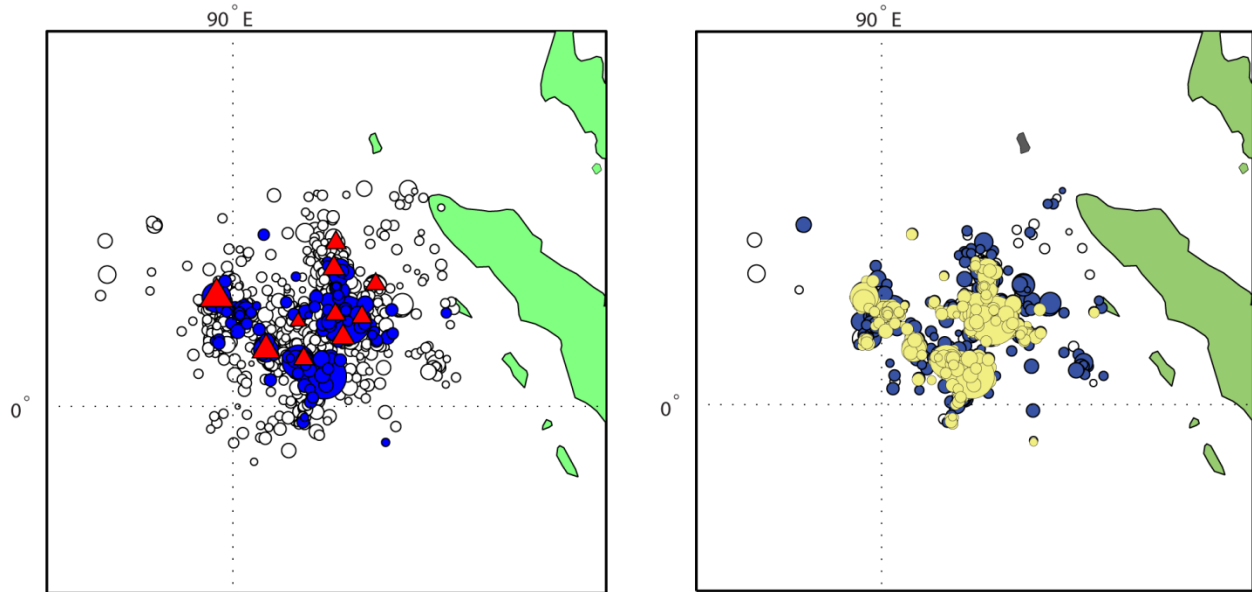


Figure 11 (Left) Red triangles are events used to create power detectors. Blue circles are 1-day power detections and white circles are catalog events. (Right) Yellow circles are subspace detections that match those of the spawning detectors. Blue circles are the additional detections over the 44-day period.

Figure 11(Right) shows the 532 subspace detections for the 44-day period (blue) atop the 4-detector power detections (white). The yellow circles are the detections in common between the 10 power detectors and the subspace detectors. The footprint of the 10 subspace detectors is larger than the area covered by the basis events, but still does not come close to covering the entire aftershock zone.

Compared to the high-rank subspace detectors of the previous section this configuration produced 68% as many detections. All but 92 of the 4-detector power detections shown in Figure 8 were reproduced. There were also an additional 88 detections. Under the assumptions used throughout this report, the analyst workload reduction factor for this set of detectors is 98% $((532-10)/532 * 100\%)$. From that standpoint, this set of detectors is very successful. Although there would have been a day's delay while the basis was being formed, for the rest of the sequence, about 43% of the events would have been detected and correctly classified by this system.

In this section we have shown that a high-rank subspace detector can be created that has a footprint larger than the domain of its basis detections and that can detect many events within its footprint that are not part of its basis. As part of a system designed to exploit these kinds of detections, it could provide a significant reduction in analyst workload. However, because the subspace templates do not evolve, the configuration is not able to track a sequence that migrates away from the region that was active during the period of basis formation.

Conclusions

We have presented three methods for creating suites of multi-rank subspace detectors and have applied them to the 2012 Sumatra earthquake sequence. In the first strategy, correlators are spawned directly from power detections and are allowed to run to the end of the sequence. Correlation clustered detections are then used to create multi-rank subspace detectors which are run against the entire sequence. In the second strategy, power detectors are run for the entire duration of the sequence. Nearly the entire set of power detections is used to create one or more subspace detectors with a high energy capture value. The final strategy was to build subspace detectors directly from un-clustered power detections, but restricted to only the first day's detections.

All three strategies required the detections from the power detectors to be filtered by slowness and by SNR. We also required the subspace templates to be 50 seconds or more in length. Without these restrictions, significant numbers of false detections could be produced by the subspace detectors as the detection threshold was decreased.

None of the detectors presented here are as sensitive as the system presented by Bobrov et al, 2012, and none could be directly plugged into the existing pipeline at the IDC. However, we think that detectors with the ability to reliably detect and classify a large number of events using data recorded by a single array could play an important role in future pipeline systems. They could allow much more robust association of detections during times of high seismicity, and could potentially provide a basis for processing data in groups, thus improving efficiency.

References

- Bobrov, D., J. Coyne, J. Given, U. Khukhuudei, I. Kitov, K. Sitnikov, S. Spiliopoulos, and L. Zerbo (2012c). Sumatera 2012 Aftershocks: Reviewed Event vs. Waveform CrossCorrelation Bulletin, Proceedings of 2012 Monitoring Research Review: GroundBased Nuclear Explosion Monitoring Technologies, Albuquerque, New Mexico, 18-20 September, 754-763.
- Bunch, J. R. and C. P. Nielsen (1978). Updating the singular value decomposition, Numerical Math., 31, pp. 111-129.
- Harris, D. (2006). Subspace detectors: Theory, Lawrence Livermore Natl. Lab. Rep. UCRL-TR-222758, 46 pp., Lawrence Livermore Natl. Lab., Livermore, California (Available at <http://www.llnl.gov/tid/lof/documents/pdf/335299.pdf>)
- Harris, D. B. and D. A. Dodge (2011). An autonomous system of grouping events in a developing aftershock sequence, Bull. Seismol. Soc. Am. 101:2, 763–774.
- Kvaerna, T., D. B. Harris, S. J. Gibbons, and D. Dodge (2011) A Matched Field Processing Framework for Coherent Detection Over Local and Regional Networks, Proceedings of 2011 Monitoring Research Review: Ground-Based Nuclear Explosion Monitoring Technologies, Albuquerque, New Mexico, 13-15 September, 292-301.
- Kvaerna, T., D. B. Harris, S. J. Gibbons, and D. Dodge (2012) Adapting Pipeline Architectures to Track Developing Aftershock Sequences and Recurrent Explosions Proceedings of 2012 Monitoring Research Review: Ground-Based Nuclear Explosion Monitoring Technologies, Albuquerque, New Mexico, 18-20 September, 776-785.